

AD-A238 925



DOCUMENTATION PAGE

Form Approved
OMB No 0704-0188

1a REPORT SECURITY CLASSIFICATION Unclassified			1b RESTRICTIVE MARKINGS	
2a SECURITY CLASSIFICATION AUTHORITY			3 DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.	
2b DECLASSIFICATION/DOWNGRADING SCHEDULE				
4 PERFORMING ORGANIZATION REPORT NUMBER(S)			5 MONITORING ORGANIZATION REPORT NUMBER(S)	
6a NAME OF PERFORMING ORGANIZATION Harvard University		6b OFFICE SYMBOL (If applicable)	7a NAME OF MONITORING ORGANIZATION Cognitive Science Program	
6c ADDRESS (City, State, and ZIP Code) 33 Kirkland Street Cambridge, MA 02138			7b ADDRESS (City, State, and ZIP Code) Office of Naval Research (Code 1142CS) 800 N. Quincy St., Arlington, VA 22217-5000	
8a NAME OF FUNDING/SPONSORING ORGANIZATION		8b OFFICE SYMBOL (If applicable)	9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-90-J-1826	
8c ADDRESS (City, State, and ZIP Code)			10 SOURCE OF FUNDING NUMBERS	
			PROGRAM ELEMENT NO 61153N	PROJECT NO RR04206
			TASK NO RR04206-	WORK UNIT 4422560-
			01	01
11 TITLE (Include Security Classification) Components of Picture Naming				
12 PERSONAL AUTHOR(S) Stephen M. Rosslyn				
13a TYPE OF REPORT Final Technical		13b TIME COVERED FROM 4/15/90 TO 4/14/91	14 DATE OF REPORT (Year, Month, Day) 91-7-15	
15 PAGE COUNT 29				
16 SUPPLEMENTARY NOTATION Reproduction in whole or in part is permitted for any purpose of the US Government.				
17 COSATI CODES			18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP		
05	08			
19 ABSTRACT (Continue on reverse if necessary and identify by block number) Additive factors methodology was used to study the componential structure of high-level visual processes. Subjects either named pictures of common objects or decided whether words named pictures; response times and error rates were recorded. The pictures were manipulated in ways that were thought to selectively affect distinct processing stages or that were thought to affect the same processing stage. For example, spurious parts were added, which should have affected matching in associative memory, and random line segments were placed over the picture, which should have affected relatively early "feature extraction" processes. As expected, manipulations that should have affected different stages produced additive effects on response times, whereas manipulations that should have affected the same stage did not produce discrete increments of additional time.				
20 DISTRIBUTION/AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS			21 ABSTRACT SECURITY CLASSIFICATION	
22a NAME OF RESPONSIBLE INDIVIDUAL Susan E. Chipman			22b TELEPHONE (Include Area Code) (202) 696-4318	22c OFFICE SYMBOL ONR 1142CS

DD Form 1473, JUN 86

Previous editions are obsolete

S/N 0102-LF-014-6603

91-05989



FINAL REPORT
Components of Picture Naming
(N00014-90-J-1826)
S. M. Kosslyn, PI

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution	
Availability Codes	
Dist	Avail and/or Special
A-1	

Kosslyn final report
-1-

Any machine that can behave as human beings do will have to be very complicated. Perhaps the only way to construct such machines, either through evolution or manufacture, relies on the principle of modularity (cf. Simon, 1981). And in fact, research in cognitive science has demonstrated repeatedly that many, if not all, cognitive faculties are accomplished by the joint action of a host of component processes. However, although it often is clear that numerous component processes must be at work, it is not often clear how best to characterize them. This research was designed to discover the component processes used during visual object recognition and identification.

The availability of new brain imaging technologies offers new opportunities for characterizing the functional architecture of cognitive systems. For example, Posner, Petersen, Fox & Raichle (1988) used positron emission tomography (PET) to study the components of information processing used during reading. However, the studies by Posner, Raichle and their colleagues (and all of the best work in this field) rely on a *subtractive methodology*. They try to isolate components by subtracting the pattern of brain activation evoked by a control task from that evoked by a test task, with the assumption that the test task recruits all of the processes used in the control task plus one additional process; hence, the subtraction is presumed to leave only the pattern of brain activation evoked by the additional component in the test task. However, the dangers of using this methodology have been clear since the time of the Wurtzburg school (circa 1913). Specifically, altering a control task to produce the test task may (or may not) lead subjects to change strategy, and use a different set of processes from those used in the control task. Thus, the subtraction may or may not implicate a specific component process. So far there has been no attempt to specify when it is appropriate to subtract one pattern of activation from another in order to isolate the contribution of a single processing component.

The research performed here served two purposes: First, it provided evidence for a specific theory of the component processes used in visual object recognition and identification. Second, it was an attempt to develop a more precise way to use the subtraction methodology to observe the pattern of brain activation evoked by a particular processing component. Thus, this work positioned us to use PET scanning to obtain convergent evidence for a set of processing components, and to study how these components are implemented in the brain.

Logic of the method

The key to the methodology developed here is to identify variables that selectively affect individual processing components. Following the logic of S. Sternberg (1969), we argue that *if processing components are in fact independent, then variables that affect a given process should do so in the same way regardless of what other processes are used to perform a task*. Sternberg illustrated this logic in a memory scanning task, in which subjects memorize a list of digits, and then decide whether a probe digit was in the memorized set. Sternberg reasoned that varying the perceptual quality of the probe should affect an encoding stage, but not a subsequent stage in which the encoded probe is compared to items on the memorized list; in contrast, varying the size of the list should affect the time to search for the probe, but should not affect encoding it. If the two stages are independent, then variables that affect the processing in one should do so in the same way, regardless of how other variables affect processing in other stages. And in fact, the two variables (probe quality and set size) had statistically independent effects. Thus, these results provided support for the independence of the processes that encode digits and search memorized lists. (Note that if such results fail to be obtained, this does not necessarily imply that the assumption is incorrect, given the possibility of cascade and interactive processes [see McClelland, 1979]; however, if such additivity is obtained, we have support for the assumption.)

All the experiments reported here followed the same logic. First, we identified variables that the theory (Kosslyn, Flynn, Amsterdam, and Wang, 1990; Kosslyn, in press) predicts would affect processing in just one subsystem. For example, according to the theory, a "preprocessing subsystem" extracts invariant properties of visual input, such as parallel lines and points of intersection (see Lowe, 1987a, b). Therefore, we inferred that superimposing randomly placed line fragments over a picture would impede, or "stress," the functioning of this subsystem. Our inference was based on the idea that the line fragments would create more invariant properties, forcing the preprocessing subsystem to expend more resources (time, blood flow, metabolism, etc.) to accomplish its function. We examined either two or three levels of the particular variable in each case. In this example, we superimposed either no noise at all, "sparse" noise, or "dense" noise over the pictures, for a total of three levels.

Second, we we orthogonally combined at least two of these variables in the design of a single experiment. The task was either picture naming or picture name verification (see Kosslyn & Chabris, 1990), and the dependent measure was mean response time.¹ In each experiment, we chose variables predicted to affect different subsystems. Thus, in any single experiment, to the extent that the theory is correct, the variables should not interact with one another (Sternberg, 1969). The theory predicts that because the processes are distinct, the increment of response time caused by stressing one process should be perfectly additive with the increment caused by stressing the other. For example, we combined the preprocessing variable (added noise fragments) with a variable thought to influence the process of matching input to previously stored memories (in the "pattern activation subsystem"), namely the number of nonaccidental properties removed from the original picture. The pattern activation subsystem posited by Kosslyn et al. (1990) operates via a process of constraint satisfaction; incoming the nonaccidental properties and their spatial relations are matched to those of objects and parts stored in visual memory. To the extent that the constraints are weakened, it will be more difficult to implicate a single stored representation.

Third, to address the problem of affirming the null hypothesis, we adopt the logic of "deductive testing of theories" (Popper, 1968), or falsificationism. Each experiment can be viewed as an attempt to falsify the theory by observing an interaction between one or more of the key variables. If the interaction is statistically significant, then there is evidence that the variables do not affect distinct subsystems. However, such a null finding could occur for one or both of two reasons: One or more of the variables do not actually affect the subsystems we inferred, or the subsystems they do affect are not distinct in the brain. If the interaction is nonsignificant, we can perform a power analysis (Cohen, 1988) to show that we had a high probability of detecting it were various effect sizes present in the population. Thus, by performing several experiments of reasonable power and consistently finding no interactions, we continually fail to falsify the theory, and therefore build support for it.

Experiment 1

This experiment tested the distinction between the unimodal visual *pattern activation* subsystem (PA) and the multimodal *associative memory* subsystem (AM). The pattern activation subsystem is thought to be based on processes in the inferior temporal lobe, and matches visual input to stored visual memories. In contrast, associative memory is thought to be distributed, but depends on processes in the superior-posterior temporal lobe. A wide variety of information is stored in associative memory, including descriptions of the arrangements of parts of an object. According to Kosslyn et al. (1990), if the initial match in the pattern activation subsystem is not very good, the output is treated as an hypothesis to be tested. Hypotheses are tested when information about the candidate object is accessed in associative memory, and used to guide top-down search for a distinctive part.

To stress the pattern activation subsystem we degraded pictures of common objects by removing parts of the objects or leaving all of the parts intact. We inferred that because the pattern activation subsystem matches visual input to stored visual memories, it should have more difficulty making a match if the input is missing some of the parts of the canonical representation of the object in question.

To stress associative memory we added extraneous parts from other objects to the pictures. We inferred that this manipulation would affect associative memory processing because the irrelevant parts would temporarily provide positive evidence for the presence of the object(s) from which they came. This would slow down the process of constraint satisfaction used to derive a single name from the visual input to associative memory (Kosslyn & Chabris, 1990; see also Kosslyn, in press).

Furthermore, we concluded that the degradation would not affect associative memory because the symbol that is the output of the pattern activation subsystem—which in turn becomes input to associative memory—would be the same once it had made a match regardless of the strength of the match (which would depend on the number of parts removed). We similarly concluded that adding parts would not affect the pattern activation subsystem because it is hypothesized to match all visual inputs to all visual memories in parallel, so if parts from more than one object were present, it would take no longer than if only parts from a single object were present.

Finally, we presented the pictures only briefly to prevent subjects from making eye movements and/or shifting their attention over the image to search for distinctive parts of the objects. This restricted the amount of information subjects had to perform top-down hypothesis testing, and hence should have amplified the effect of having to do so.

Method

Subjects. Thirty two high school or college students volunteered to participate as paid subjects. Half were male and half were female. None had previously participated in any of the other experiments reported here, nor had they ever seen any of the pictures used in this experiment.

Materials. Stimuli were black-and-white line drawings of 48 common objects from the set described by Snodgrass and Vanderwart (1980), as digitized by Brooks (1985). Only pictures with name agreement greater than 80% and relatively high familiarity ratings were included, and the number of animals was deliberately minimized. Appendix 1 lists the objects used.

Bitmapped images of each picture were modified to produce “degraded” versions of each; this was done by removing one or more distinguishable parts from the objects. For example, the picture of a football was degraded by removing the laces and one of the stripes. Then, we added one extraneous part to each object by modifying both the intact and the degraded versions of the picture. In the case of the football, the head of a screw was placed next to it as though it were sticking out of the side of the ball. Thus we created four versions of each picture: original, degraded, original/added, and degraded/added, for a total of 192 possible stimuli. Examples of the stimuli are shown in Figure 1.

Insert Figure 1 About Here

Procedure

Each subject completed 48 trials, seeing 12 different objects in each of the four conditions, and no object in more than one condition. This was done to ensure that previously viewing an object in one condition could not influence a subsequent response to the same object in a different condition. The allocation of pictures to conditions was varied systematically so that within a

counterbalancing group of four subjects, each picture appeared exactly once in each condition. One pseudorandom trial order was created, subject to the constraint that no more than three consecutive trials could include the same type of picture (nondegraded, degraded, nonadded, added). Three more trial orders were created by reversing the original order, swapping the first and second halves of the original order, and reversing this new swapped order, for a total of four trial orders. Two subjects were tested for each combination of counterbalancing group and trial order (one male and one female), for a total of 32 subjects.

At the beginning of each trial the screen was blank (white) for 2000 ms. Then a small asterisk appeared in the center of the screen, where subjects were told to fixate their attention throughout the experiment. After 500 ms the asterisk was replaced by the stimulus picture, which remained for 100 ms. After the picture disappeared the screen remained blank until subjects responded, after which the next trial began immediately. Subjects were instructed to speak the most appropriate name for the "major" object in each stimulus as quickly and accurately as they could after its appearance. An experimenter recorded the actual words spoken by the subjects and the computer recorded response time.

The MacLab program (Costin, 1988) running on an Apple Macintosh Plus computer with a Polaroid CP-50 anti-glare filter was used to present stimuli and record response times. A Radio Shack model 33-992C microphone was connected to a Lafayette Instrument model 18010 voice-activated relay, which was connected to the computer keyboard. This configuration allowed the MacLab program to record subject response times with millisecond accuracy.

Subjects were tested individually in sessions of approximately 15 minutes. After filling out a consent form, subjects read the instructions and completed eight practice trials (two in each condition). This practice session was identical for all subjects, and used objects that did not appear in the experimental trials. Before completing the experimental trials, subjects were shown copies of the stimuli they had seen in the practice session. At the end of the session they filled out a debriefing sheet and a short version of the Edinburgh Handedness Inventory (Oldfield, 1971).

Results

An analysis of variance was conducted on mean response times for each condition for each subject. Degradation (whether or not parts were removed) and addition (whether or not parts were added) were within-subject factors and subject was the random effect. Before computing the means, we removed from the raw data times from incorrect responses. We then iteratively trimmed away all response times outside three standard deviations above or below the mean of the remaining response times in each cell. These procedures served to exclude response times from different distributions and outliers, thereby making the resulting distribution more normal and more reflective of the actual processing needed to perform the task.

Insert Figure 2 About Here

Figure 2 illustrates the results. As predicted, the interaction of degradation and addition was not significant, $F(1,31) = 2.32$, $p = .1375$, $r = .2641$. That is, the effect of adding parts was not significantly greater when the picture was degraded (104 ms) than when it was not degraded (63 ms). A power analysis revealed that we had at least approximately .17 probability of detecting an interaction of the effect size observed (or larger), if it in fact existed. However, we had higher power to detect larger effects were they present in the population; for instance, we had at least a .64 probability of detecting an effect of $r = .50$ or larger.

In addition, degraded pictures took longer to name than nondegraded pictures (862 ms versus 766 ms), $F(1, 31) = 50.15$, $p = .0001$, and pictures with added parts took longer to name than pictures without added parts (855 ms versus 772 ms), $F(1, 31) = 35.31$, $p = .0001$. The

results were virtually identical when error rates were analyzed; there was no hint of a speed-accuracy trade off.

Discussion

These results were as expected if nonaccidental properties are in fact extracted by a different subsystem than the one that matches input to stored representations. We found that a variable that should affect the preprocessing subsystem had effects that added onto those of a variable that should affect the pattern activation subsystem.

Experiment 2

In this experiment we again tested the pattern activation/associative memory distinction, but here the pictures remained in free view until the subjects responded. This allowed them to engage top-down processing mechanisms. We expected that if the distinction were robust, it should manifest itself under these conditions as well.

Method

Subjects. Sixteen high school or college students volunteered to participate as paid subjects, again being equally divided between the two genders. None participated in any of the other experiments reported here, nor had they ever seen any of the pictures used in this experiment.

Materials. The materials were identical to those used in Experiment 1.

Procedure

The procedure was identical to that of Experiment 1, with one exception. Each stimulus picture remained on the screen until subjects spoke their response into the microphone. (The instructions given to the subjects were modified accordingly.)

Results

Mean response times were computed and analyzed as in Experiment 1. Figure 3 illustrates the results. Again, the interaction between degradation and addition was not significant, $F(1, 15) = .03$, $p = .8566$, $r = .0474$. Here, the effect of adding parts was virtually equal when the picture was degraded (95 ms) to when it was not degraded (87 ms). A power analysis revealed that we had approximately .03 probability of detecting such a small interaction if it existed in the population, but as discussed previously, we had considerable power to detect larger effects. In addition, degraded pictures took significantly longer to name than nondegraded pictures (1003 ms versus 927 ms), $F(1, 15) = 23.67$, $p = .0002$, and pictures with added parts took significantly longer to name than pictures without added parts (1011 ms versus 920 ms), $F(1, 15) = 9.79$, $p = .0069$. The results were similar when error rates were analyzed, although the main effect of degradation was not significant, $F < 1$, and the main effect of addition was slightly reversed. However, the difference of 0.52% (9.63% errors on pictures without added parts, 9.11% errors on pictures with added parts) was not significant, $F < 1$, belying a speed-accuracy tradeoff. Thus, these results also failed to falsify the hypothesis that the pattern activation and associative memory subsystems are distinct.

Insert Figure 3 About Here

Discussion

These findings provide additional support for another distinction made by the Kosslyn et al. (1990) theory, namely between a visual memory in which more than one representation can be activated and an associative memory in which only one representation is activated by an input.

Experiment 3

With their identical results, both Experiment 1 and Experiment 2 provided support for the distinction between pattern activation and associative memory processing. We hypothesized that manipulating the exposure time of the pictures between the two experiments had no effect on the additivity of the degradation and addition variables because exposure time was influencing a third process. Kosslyn et al. (1990) posit a "visual buffer" (VB), which is an "active structure" that includes the low-level areas of the visual system and the low-level processes that act there to transform visual input. The ability of these bottom-up processes to successfully organize input for further processing could be impaired by presenting the input for a short period of time, such as the 100 ms used in Experiment 1, compared to the several hundred milliseconds consumed by the subjects in Experiment 2. To test the distinction between the visual buffer, pattern activation, and associative memory subsystems, we incorporated all three variables (exposure time, degradation, and addition) into a single experiment.

First, however, we combined the data from Experiments 1 and 2 into a single three-way ANOVA, with experiment, corresponding to exposure time (100 ms versus free view), as a new between-subjects factor. In this analysis, none of the interactions were significant, $F < 1$, $p > .30$, $r < .15$, and power less than .20 in all cases. (The main effects were all significant, as expected.) However, this could be due to the generally low power of the between-subjects design. Furthermore, the interaction between exposure time and addition was significant in the error rate analysis, $F(1,46) = 9.23$, $p = .0039$. This was due to the fact that—for some unknown reason—adding parts produced an increase of 7.68% errors in the short exposure experiment but a decrease of 0.52% errors in the long exposure experiment. In the following experiment we used a within-subjects design to explore the effects of these three variables in more detail.

Method

Subjects. Thirty two college students or individuals recruited through a newspaper advertisement, ranging in age from 17 to 34 (mean 20), volunteered to be paid subjects; as usual half were male and half were female. None had previously participated in any of the other experiments reported here, nor had they ever seen any of the pictures used in this experiment.

Materials. The materials were identical to those used in Experiment 1.

Procedure

The procedure was identical to that of Experiment 1, with the following exceptions. Each picture appeared with two exposure times, short (50 ms) and long (500 ms). Thus, each picture now appeared in eight conditions (two levels of degradation, two levels of addition, and two exposure times), which were orthogonally combined. Each subject saw six different pictures in each condition rather than twelve, and the counterbalancing groups therefore included eight subjects rather than four. The same trial orders were used, so one subject was tested for each combination of counterbalancing group and trial order, for a total of 32 subjects.

Results

Mean response times were computed and analyzed as in Experiment 1, with the addition of exposure time as a within-subject factor. As expected, none of the interactions were significant ($p > .20$, $r < .25$, and power less than .25 in all four cases), but all three main effects were significant. Short-exposure pictures took longer to name than long-exposure pictures (940 ms versus 571 ms), $F(1,31) = 570.00$, $p = .0001$; degraded pictures took longer to name than nondegraded pictures (799 ms versus 712 ms), $F(1,31) = 30.50$, $p = .0001$; and pictures with added parts took longer to name than pictures without added parts (793 ms versus 719 ms), $F(1,31) = 26.80$, $p = .0001$. The pattern of results was identical in the error rate analysis, so there were no speed/accuracy tradeoffs.

Discussion

The results of this experiment provided further support for the distinction among the three subsystems (visual buffer, pattern activation, and associative memory). Variables that should have affected each subsystem did so independently.

Experiment 4

This experiment was intended to obtain convergent evidence for our previous characterization of the pattern activation subsystem. We now selected two variables that should affect processing in the pattern activation subsystem. The design of this experiment was identical to that of Experiment 4, except that in place of the sparse grid conditions, we rotated the picture 45 degrees clockwise, and in place of the dense grid conditions, we rotated the picture 135 degrees clockwise. Jolicoeur (1985) found that naming time increases with angle of rotation from the upright position of the picture. He interprets this finding to indicate that people "mentally rotate" the pictures, although he acknowledges that the slope is an order of magnitude smaller than that typically found in mental rotation experiments. We (see Kosslyn & Chabris, 1990) interpret this finding as indicating that a rotated stimulus is matched less effectively to stored memories, which requires subsequent processing to search for distinctive parts.

If so, then rotating the figure should interact with degrading it; both variables affect the ease of matching input to stored visual representations.

Method

Subjects. Twenty individuals recruited through a newspaper advertisement, ranging in age from 17 to 26 (mean 22) years, volunteered to be paid subjects; 3 were male and 17 were female. All reported being right-handed. None had previously participated in any of the other experiments reported here, nor had they ever seen any of the pictures used in this experiment.

Materials. The materials were identical to those used in Experiment 4, with the following exceptions. The noise conditions were replaced by conditions in which the picture was rotated clockwise 45 degrees (instead of superimposing a sparse noise grid) and 135 degrees (instead of superimposing a dense noise grid). Thus there were five versions of each picture: original, 30% degraded/45 degrees rotated, 30%/135, 60%/45, and 60%/135, for a total of 250 possible picture stimuli. Examples of the picture stimuli are shown in Figure 4.

Insert Figure 4 About Here

Procedure

The procedure was identical to that of Experiment 4.

Results

Mean response times were computed and analyzed as in Experiment 4, with the factor rotation (45 versus 135 degrees) replacing noise. Figure 5 illustrates the results. The interaction of degradation and rotation was not significant, $F(1,19) = 1.99$, $p = .1748$, $r = .3077$. That is, the effect of increasing rotation was not significantly greater when the picture was 60% degraded (195 ms) than when it was 30% degraded (75 ms). The power of this experiment was only .2639, and the relatively large effect size suggests that the interaction could be present in the population. And in fact, *t*-tests comparing the means for the four conditions revealed, after adjustment for four comparisons by the Bonferroni procedure, that the effect of rotation was significant for the 60% degraded pictures, $t(19) = 3.15$, adjusted $p = .02$, but was not significant for the 30% degraded pictures, $t(19) = 1.95$, adjusted $p = .2644$.

Insert Figure 5 About Here

The interaction between rotation and correct response was significant, $F(1,19) = 9.55$, $p = .0060$, as was the interaction between degradation and correct response, $F(1,19) = 17.00$, $p = .0006$. The three-way interaction between degradation, rotation, and correct response was not significant, $F(1,19) = 1.54$, $p = .2304$. As in Experiment 4, subjects responded faster to 30% than to 60% degraded pictures (877 ms versus 1080 ms), $F(1,19) = 17.70$, $p = .0005$, and also responded faster to 45 degree than to 135 degree rotated pictures (911 ms versus 1046 ms), $F(1,19) = 12.52$, $p = .0022$. "Yes" responses were faster than "no" responses (941 ms versus 1016 ms), $F(1,19) = 7.93$, $p = .0110$. However, as in Experiment 4, the effect was reversed in the error rate analysis, with "yes" responses resulting in more errors than "no" responses (4.50% versus 2.25%), $F(1,19) = 4.54$, $p = .0464$. There were no other results of interest in the error rate analysis.

Discussion

As expected, we found that subjects required longer to identify highly degraded pictures when they were rotated, but did not require more time to identify barely degraded pictures when they were rotated. However, although these results were in the expected direction, the failure to find a significant interaction leads us to interpret them with caution.

Experiment 5

In this experiment sought evident that the pattern activation subsystem is distinct from a different subsystem, the "preprocessing" subsystem. According to Kosslyn (in press; see also Kosslyn et al., 1990) this subsystem receives input from the visual buffer and identifies and marks the "nonaccidental properties" (Lowe, 1987a, 1987b; see also Biederman, 1987) on the image. These nonaccidental properties include parallel lines, points of intersection, and other aspects of an object's image that remain the same regardless of the configuration of the object or the viewpoint of the observer. This information presumably is then provided to the pattern activation subsystem along with the original image for comparison to stored visual memories. Note that the function of the preprocessing subsystem is mainly stimulus-driven, since it does not store specific visual memories, but it can be "tuned" with top-down, context-specific information acquired through perceptual learning (e.g. Biederman & Shiffrar, 1987).

To stress the preprocessing subsystem we superimposed pseudorandom patterns of lines, a sort of "visual noise," over pictures of common objects. We inferred that since the preprocessing subsystem compulsively encodes all nonaccidental properties in its input image, the coincidental intersections, parallel lines, and other features produced by the noise would increase the amount of processing it had to do, thus taxing its abilities.

To stress the pattern activation subsystem we once again degraded the pictures, but this time we did so by removing portions of the pictures that would contribute to the formation of nonaccidental properties. For example, we removed segments of parallel lines, or we removed one of the two lines in an intersecting pair. We inferred that this manipulation would have an effect similar to that of removing parts in Experiments 1-3. As was noted earlier, when constraints are eliminated, the matching process is less effective.

Furthermore, we assumed that the noise would not impede pattern activation processing because once the nonaccidental properties have been extracted, the matching can proceed as normal (following the same logic as Sternberg's reasoning about the additivity of encoding and comparison stages). We also concluded that the degradation would not impede the preprocessing subsystem because the removed components of nonaccidental properties would, if anything, *reduce* the processing load.

Finally, in this experiment we modified the task. Instead of picture *naming*, we used picture *name verification*. Subjects saw the picture, heard a name, and had to decide whether the name was an appropriate one for the picture.

Method

Subjects. Twenty high school students, college students, or university employees, volunteered to participate as paid subjects; half were male and half were female. All reported being right-handed. None had previously participated in any of the other experiments reported here, nor had they ever seen any of the pictures used in this experiment.

Materials. Stimuli were black-and-white line drawings of 50 common objects from the set described by Snodgrass and Vanderwart (1980). Here we used higher quality digitized versions supplied directly by Snodgrass (personal communication). Only pictures with name agreement greater than 80% and relatively high familiarity ratings and word frequencies were included, and the number of animals was reduced further below that in Experiments 1–3. Twenty five of the objects were classified as “straight” by the experimenters based on the type of line segments that predominated in their pictures; the remaining 25 objects were considered “curvy.”

Bitmapped images of each picture were modified to produce two “degraded” versions of each. The first degraded version was made by removing nonaccidental properties from the image until approximately 30% of the original pixels (or bits in the bitmap) were gone. The second degraded version was produced by continuing the process until approximately 60% of the original pixels were gone. Then, we superimposed two different grids over each degraded picture: a “sparse” and a “dense” grid, as determined by the number and density of the lines in the grids. Straight pictures received only grids of curvy lines, and curvy pictures received only grids of straight lines.² The straight grids contained slightly more lines than the corresponding curvy grids, since the number of lines in each grid was adjusted in pilot studies to ensure that both sparse grids appeared equally noisy and both dense grids appeared equally noisy. In addition, to ensure that subjects would not become accustomed to the appearance of the grids and begin to ignore them, we created four variants of each grid (density and line type) by rotating the originals 0, 90, 180, or 270 degrees clockwise. Each object appeared with one of the four possible grids for its condition. Finally, we included stimuli in which the original picture appeared, nondegraded and with no grid. Thus we created five versions of each picture: original, 30% degraded/sparse grid, 30%/dense, 60%/sparse, and 60%/dense, for a total of 250 possible picture stimuli. Examples of the picture stimuli are shown in Figure 6.

Insert Figure 6 About Here

Each object was then paired with two words, the correct name and an incorrect name for the object. The incorrect names were correct names of other objects in the experiment. Appendix 3 lists the objects used, as well as the incorrect names paired with each object; note that the incorrect names for straight objects were correct names of curvy objects, and vice-versa. Thus, since each of the 250 pictures could occur as a “yes” trial and as a “no” trial, there were 500 total possible trials.

The name of each picture was recorded using the Farallon Computing MacRecorder sound digitizer and the SoundEdit program on a Macintosh II computer. Sound was sampled with eight-bit resolution at a frequency of 11 KHz.

Procedure

Each subject completed 50 trials, seeing 10 different objects in each of the five conditions, and no object in more than one condition. As in Experiments 1–3, this was done to ensure that previously viewing an object in one condition could not influence a subsequent response to the same object in a different condition. Of the ten objects in each condition, five were in “yes” trials

and five were in "no" trials. The allocation of pictures and sounds to conditions was varied systematically so that within a counterbalancing group of ten subjects, each picture appeared exactly once in each condition as a "yes" trial and once in each condition as a "no" trial. One pseudorandom trial order was created, subject to the constraint that no more than three consecutive trials could include the same type of picture (original, 30% degraded, 60% degraded, sparse noise, dense noise) or the same type of response (yes, no). Two subjects were tested for each counterbalancing group, for a total of 20 subjects.

At the beginning of each trial a large exclamation point appeared in the center of the screen, where subjects were told to fixate their attention throughout the experiment. Subjects were instructed to press the space bar when they were ready to proceed with the trial. At that point the screen became blank for 500 ms. Then the picture appeared and remained on the screen for 500 ms, after which the sound was played. The next trial began as soon as the subject responded. Subjects were instructed to press the appropriate key—"yes" with their index finger or "no" with their middle finger of their right hand—as quickly and accurately as possible after they heard the sound. (Response time was measured from the onset of the sound, not the picture.)

The apparatus was the same as in Experiments 1–3, but the voice-activated relay and microphone were not used, and the MacLab program recorded both responses and times. Subjects were tested individually in sessions of approximately 15 minutes. After filling out a consent form, subjects read the instructions and completed ten practice trials (one "yes" trial and one "no" trial in each condition). This practice session was identical for all subjects, and used objects and sounds that did not appear in the experimental trials. At the end of the session they filled out a debriefing sheet and a short version of the Edinburgh Handedness Inventory (Oldfield, 1971), which was later examined to ensure that all subjects were right-handed.

Results

Mean response times were computed as in Experiment 1. We then conducted an ANOVA that included degradation (30% or 60%), noise (sparse or dense), and correct response type (yes or no) as within-subject factors but excluded the nondegraded, noiseless condition. Figure 7 illustrates the results. As predicted, the interaction of degradation and noise was not significant, $F(1,19) = 1.17$, $p = .2921$, $r = .2413$. That is, the effect of increasing noise density was not significantly greater when the picture was 60% degraded (119 ms) than when it was 30% degraded (43 ms). A power analysis revealed, however, that we would have detected such a small interaction only about 19% of the time with 20 subjects in this experiment. Furthermore, t-tests comparing the means for the four conditions revealed, after adjustment for four comparisons by the Bonferroni procedure, that the effect of degradation was significant for the pictures with sparse grids, $t(19) = 4.20$, adjusted $p = .002$, as well as for the pictures with dense grids, $t(19) = 3.70$, adjusted $p = .006$. However, the effect of noise was not significant for either 30% or 60% degraded pictures, $t(19) = 1.65$, adjusted $p = .46$ and $t(19) = 1.84$, adjusted $p = .32$ for the two comparisons respectively.

Insert Figure 7 About Here

None of the other interactions were significant ($p > .25$ in all three cases), but subjects evaluated 30% degraded pictures faster than to 60% degraded pictures (802 ms versus 1048 ms), $F(1,19) = 20.49$, $p = .0002$, and pictures covered by sparse noise faster than pictures covered by dense noise (885 ms versus 965 ms), $F(1,19) = 5.48$, $p = .0303$. "Yes" responses were faster than "no" responses (888 ms versus 962 ms) but the effect only approached significance, $F(1,19) = 4.22$, $p = .0540$. However, the effect was reversed in the error rate analysis, with "yes" responses resulting in more errors than "no" responses (8.88% versus 2.00%), $F(1,19) = 9.95$, $p = .0052$. There was also an interaction between degradation and response type in the error rates, $F(1,19) = 5.40$, $p = .0314$. Otherwise, the pattern of results was identical in the error rate analysis.

Discussion

Although the results were as expected, the lack of power makes us cautious in interpreting them. In general, we must be very cautious about accepting even predicted null findings when the power is low.

General Discussion

The results of these experiments generally support the functional decomposition posited by Kosslyn et al. We found evidence for the distinction between the visual buffer, preprocessing, pattern activation, and associative memory/top-down processing subsystems.

However, we became increasingly leary of interpreting null findings as strong evidence for the functional distinctions posited by the theory. In most cases, even with sizable n 's, our power was relatively low. The logic of additive factors inherently relies on affirmation of null effects, which may be an insurmountable limitation.

More compelling evidence for the functional distinctions will be forthcoming in the next stage of this research, when we administer some of these tasks while subjects are undergoing PET scanning. If we find that manipulations that stress different subsystems evoke more blood flow in different brain areas, this will be strong evidence for the decomposition.

Footnotes

1. Analyses of error rates were also examined; however, error rates are less relevant for our purposes because they lack the additivity properties of mean response times. See Sternberg (1969) for further discussion of this issue.
2. The grids were dense and composed of thick lines, so the use of "contrasting" grid lines served to attenuate the obscuring effect, which may have been overwhelming otherwise. However, in an experiment present in progress we orthogonally combined straight and curvy grids with straight and curvy pictures to examine this possibility more closely.

References

- Brooks, J. O. III. (1985). Pictorial stimuli for the Macintosh computer. *Behavior Research Methods, Instruments, and Computers*, 17, 409-410.
- Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences*. Hillsdale, NJ: Erlbaum.
- Costin, D. (1988). MacLab: A Macintosh system for psychology labs. *Behavior Research Methods, Instruments, and Computers*, 20, 197-200.
- Jolicoeur, P. (1985). The time to name disoriented natural objects. *Memory and Cognition*, 13, 289-303.
- Jolicoeur, P., Gluck, M., & Kosslyn, S. M. (1984). Pictures and names: Making the connection. *Cognitive Psychology*, 16, 243-275.
- Kosslyn, S. M. (in press). A cognitive neuroscience of visual cognition: Further developments. In R. Logie & M. Denis (Eds.), *Advances in Imagery Research*. Hillsdale, NJ: Erlbaum.
- Kosslyn, S. M., & Chabris, C. F. (1990). Naming Pictures. *Journal of Visual Languages and Computing*, 1, 77-95.
- Kosslyn, S. M., & Koenig, O. (in press). *Wet Mind: The New Cognitive Neuroscience*. New York: Free Press.
- Kosslyn, S. M., Flynn, R. A., Amsterdam, J. B., & Wang, G. (1990). Components of high-level vision: A cognitive neuroscience analysis and accounts of neurological syndromes. *Cognition*, 34, 203-277.
- Lowe, D. G. (1987a). Three-dimensional object recognition from single two-dimensional images. *Artificial Intelligence*, 31, 355-395.
- Lowe, D. G. (1987b). The viewpoint consistency constraint. *International Journal of Computer Vision*, 1, 57-72.
- Kucera, H., & Francis, W. N. (1967). *Computational Analysis of Present-day American English*. Providence, RI: Brown University Press.
- McClelland, J. L. (1979). On the time-relations of mental processes: An examination of systems of processes in cascade. *Psychological Review*, 86, 287-330.
- Oldfield, R. C. (1971). The assessment and analysis of handedness: The Edinburgh Inventory. *Neuropsychologia*, 9, 97-113.
- Popper, K. R. (1968). *The Logic of Scientific Discovery*, 2nd edition. New York: Harper & Row.
- Posner, M. I., Petersen, S. E., Fox, P. T., & Raichle, M. E. (1988). Localization of cognitive operations in the human brain. *Science*, 240, 1627-1631.
- Simon, H. A. (1981). *The Sciences of the Artificial*, 2nd edition. Cambridge, MA: MIT Press.
- Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 174-215.
- Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. *Acta Psychologica*, 30, 276-315.

Appendix 1
Objects used in Experiments 1-3

anchor
bicycle
bird
bus
butterfly
car
cat
chair
clock
cow
dog
duck
elephant
fish
football
frog
glove
goat
helicopter
horse
house
iron
kite
lion
lobster
monkey
motorcycle
ostrich
owl
piano
pig
pitcher
rabbit
rhinoceros
rocking chair
shirt
shoe
snail
spider
squirrel
swan
telephone
toaster
train
truck
violin
watch
well

Appendix 2
Power of *t*-test of $r = 0$ at $\alpha = .05$, two-tailed
(extracted from Cohen, 1988, p.92)

<i>n</i>	r_c	.20	.30	.40	.50
16	.497	.11	.21	.35	.53
20	.444	.14	.25	.43	.64
32	.349	.20	.39	.64	.85
48	.285	.28	.55	.82	.96

Notes: r_c is the effect size needed for significance at the .05 level (two-tailed); numbers in the four *r* columns are the probabilities of observing a significant effect of the given size or larger assuming it is present in the population.

Appendix 3
Objects and incorrect names used in Experiments 4-5

Straight Objects

bed	(watch)
book	(violin)
broom	(umbrella)
bus	(tree)
car	(telephone)
chair	(snowman)
comb	(shoe)
desk	(sandwich)
helicopter	(sailboat)
house	(rabbit)
ladder	(pumpkin)
lock	(pot)
motorcycle	(mushroom)
pants	(leaf)
piano	(kangaroo)
refrigerator	(iron)
ruler	(hand)
saltshaker	(grapes)
shirt	(flower)
scissors	(drum)
stool	(chain)
table	(cake)
toaster	(butterfly)
train	(belt)
truck	(bell)

Curvy Objects

bell	(truck)
belt	(train)
butterfly	(toaster)
cake	(table)
chain	(stool)
drum	(scissors)
flower	(shirt)
grapes	(saltshaker)
hand	(ruler)
iron	(refrigerator)
kangaroo	(piano)
leaf	(pants)
mushroom	(motorcycle)
pot	(lock)
pumpkin	(ladder)
rabbit	(house)
sailboat	(helicopter)
sandwich	(desk)
shoe	(comb)
snowman	(chair)
telephone	(car)
tree	(bus)
umbrella	(broom)
violin	(book)

watch (bed)

Appendix 4
Curviness ratings for Snodgrass and Vanderwart (1980) pictures
(1 = least curvy, most straight; 7 = most curvy, least straight)

12 college students, graduate students, and research assistants judged the curviness of each of the 260 pictures in the Snodgrass and Vanderwart (1980) set on a scale of 1 to 7, with 1 being the most straight and 7 being the most curvy. We then converted each judge's ratings into z-scores and used the average z-score across the 12 judges as the measure of curviness. The average interjudge correlation was $r = .6317$, and the Spearman-Brown reliability was $R = .9537$.


#	Name	Rating	z-score
1	accordion	0.93	-1.13
2	airplane	1.51	-0.75
3	alligator	3.66	0.67
4	anchor	2.29	-0.23
5	ant	3.33	0.45
6	apple	4.45	1.19
7	arm	2.63	-0.01
8	arrow	0.00	-1.75
9	artichoke	4.63	1.31
10	ashtray	2.06	-0.39
11	asparagus	1.36	-0.85
12	axe	0.97	-1.11
13	baby carriage	3.17	0.34
14	ball	4.82	1.44
15	balloon	4.40	1.16
16	banana	3.23	0.38
17	barn	1.79	-0.57
18	barrel	3.05	0.27
19	baseball bat	1.27	-0.91
20	basket	1.22	-0.95
21	bear	3.83	0.78
22	bed	1.45	-0.79
23	bee	3.48	0.55
24	beetle	3.29	0.42
25	bell	3.57	0.61
26	belt	3.51	0.57
27	bicycle	3.29	0.42
28	bird	3.77	0.74
29	blouse	2.78	0.09
30	book	0.31	-1.55
31	boot	2.67	0.01
32	bottle	2.64	0.00
33	bow	3.80	0.76
34	bowl	3.84	0.79
35	box	0.23	-1.60
36	bread	2.43	-0.14
37	broom	1.08	-1.04
38	brush	1.88	-0.51
39	bus	1.64	-0.67
40	butterfly	3.58	0.62
41	button	4.56	1.26
42	cake	4.01	0.90
43	camel	4.04	0.92
44	candle	1.94	-0.47

45	cannon	2.56	-0.06
46	cap	3.56	0.60
47	car	1.57	-0.71
48	carrot	2.20	-0.30
49	cat	3.70	0.70
50	caterpillar	2.14	-0.33
51	celery	2.84	0.13
52	chain	2.72	0.05
53	chair	0.88	-1.17
54	cherry	3.71	0.70
55	chicken	3.70	0.70
56	chisel	1.12	-1.01
57	church	0.81	-1.22
58	cigar	1.27	-0.91
59	cigarette	1.17	-0.98
60	clock	2.62	-0.02
61	clothespin	0.71	-1.28
62	cloud	4.41	1.16
63	clown	4.60	1.29
64	coat	1.55	-0.73
65	comb	0.70	-1.29
66	corn	2.82	0.11
67	couch	0.94	-1.13
68	cow	4.02	0.90
69	crown	3.51	0.57
70	cup	3.84	0.79
71	deer	3.37	0.48
72	desk	0.47	-1.44
73	dog	3.36	0.47
74	doll	3.38	0.49
75	donkey	3.49	0.56
76	door	0.57	-1.38
77	doorknob	4.08	0.94
78	dress	2.97	0.21
79	dresser	1.11	-1.02
80	drum	3.13	0.32
81	duck	3.73	0.72
82	eagle	4.01	0.90
83	ear	4.44	1.18
84	elephant	3.73	0.72
85	envelope	0.69	-1.30
86	eye	4.15	0.99
87	fence	0.50	-1.42
88	finger	2.13	-0.34
89	fish	3.61	0.64
90	flag	1.73	-0.61
91	flower	3.99	0.89
92	flute	0.51	-1.42
93	fly	2.87	0.15
94	foot	3.09	0.29
95	football	3.97	0.87
96	football helmet	4.28	1.08
97	fork	2.01	-0.42
98	fox	3.40	0.50
99	french horn	4.96	1.53


100	frog	4.33	1.11
101	frying pan	3.51	0.57
102	garbage can	3.31	0.44
103	giraffe	2.83	0.12
104	glass	2.26	-0.26
105	glasses	2.27	-0.25
106	glove	2.97	0.21
107	goat	3.44	0.52
108	gorilla	3.84	0.79
109	grapes	5.00	1.56
110	grasshopper	2.69	0.03
111	guitar	2.84	0.13
112	gun	2.34	-0.20
113	hair	2.89	0.16
114	hammer	1.31	-0.88
115	hand	3.04	0.26
116	hanger	1.98	-0.44
117	harp	2.84	0.13
118	hat	2.49	-0.11
119	heart	4.39	1.15
120	helicopter	2.58	-0.05
121	horse	3.05	0.26
122	house	0.79	-1.23
123	iron	2.54	-0.07
124	ironing board	1.02	-1.08
125	jacket	2.13	-0.34
126	kangaroo	3.53	0.58
127	kettle	2.05	-0.40
128	key	2.73	0.06
129	kite	1.52	-0.75
130	knife	1.01	-1.08
131	ladder	0.31	-1.55
132	lamp	2.41	-0.16
133	leaf	3.74	0.72
134	leg	2.70	0.04
135	lemon	4.19	1.02
136	leopard	3.69	0.69
137	lettuce	4.26	1.07
139	light switch	0.38	-1.50
138	lightbulb	3.78	0.75
140	lion	3.57	0.61
141	lips	3.80	0.76
142	lobster	3.47	0.54
143	lock	2.29	-0.24
144	mitten	3.42	0.51
145	monkey	4.12	0.97
146	moon	4.13	0.98
147	motorcycle	3.76	0.74
148	mountain	1.47	-0.78
149	mouse	3.37	0.48
150	mushroom	3.64	0.66
151	nail	0.44	-1.46
152	nail file	0.37	-1.51
153	necklace	3.35	0.46
154	needle	0.18	-1.63

155	nose	2.28	-0.24
156	nut	1.69	-0.64
157	onion	4.20	1.03
158	orange	4.54	1.25
159	ostrich	4.05	0.93
160	owl	3.69	0.69
161	paintbrush	0.77	-1.24
162	pants	1.08	-1.04
163	peach	4.18	1.01
164	peacock	4.35	1.12
165	peanut	4.38	1.14
166	pear	4.13	0.98
167	pen	0.96	-1.12
168	pencil	0.31	-1.55
169	penguin	2.69	0.03
170	pepper	3.52	0.58
171	piano	2.80	0.10
172	pig	3.65	0.66
173	pineapple	3.88	0.82
174	pipe	1.61	-0.68
175	pitcher	3.07	0.28
176	pliers	2.25	-0.26
177	plug	1.88	-0.51
178	pocketbook	2.83	0.12
179	pot	2.65	0.00
180	potato	3.49	0.56
181	pumpkin	4.25	1.06
182	rabbit	3.86	0.80
183	raccoon	3.61	0.64
184	record player	1.55	-0.73
185	refrigerator	0.43	-1.47
186	rhinoceros	3.29	0.42
187	ring	4.40	1.16
188	rocking chair	2.56	-0.06
189	roller skate	2.99	0.23
190	rolling pin	1.74	-0.60
191	rooster	3.30	0.43
192	ruler	0.08	-1.70
193	sailboat	2.77	0.08
194	saltshaker	1.80	-0.56
195	sandwich	2.35	-0.20
196	saw	0.87	-1.18
197	scissors	1.29	-0.90
198	screw	1.36	-0.86
199	screwdriver	0.99	-1.10
200	sea horse	2.73	0.05
201	seal	3.56	0.60
202	sheep	3.61	0.63
203	shirt	2.46	-0.12
204	shoe	2.61	-0.03
205	skirt	1.68	-0.64
206	skunk	3.42	0.51
207	sled	1.35	-0.86
208	snail	3.88	0.82
209	snake	4.32	1.11

210	snowman	3.87	0.81
211	sock	2.89	0.16
212	spider	2.68	0.02
213	spinning wheel	2.27	-0.25
214	spool of thread	2.69	0.03
215	spoon	2.50	-0.10
216	squirrel	3.38	0.48
217	star	0.13	-1.67
218	stool	1.78	-0.58
219	stove	0.85	-1.19
220	strawberry	3.47	0.54
221	suitcase	1.62	-0.68
222	sun	3.17	0.34
223	swan	3.81	0.77
224	sweater	2.40	-0.16
225	swing	0.09	-1.69
226	table	0.07	-1.71
227	telephone	2.76	0.08
228	television	1.46	-0.79
229	tennis racket	1.92	-0.48
230	thimble	2.25	-0.26
231	thumb	2.47	-0.12
232	tie	2.01	-0.42
233	tiger	3.09	0.29
234	toaster	0.76	-1.25
235	toe	2.44	-0.14
236	tomato	4.25	1.06
237	toothbrush	0.59	-1.37
238	top	3.62	0.65
239	traffic light	2.62	-0.02
240	train	2.10	-0.36
241	tree	4.11	0.97
242	truck	1.19	-0.97
243	trumpet	3.34	0.46
244	turtle	3.59	0.62
245	umbrella	3.34	0.45
246	vase	4.19	1.02
247	vest	2.27	-0.25
248	violin	3.27	0.41
249	wagon	1.70	-0.62
250	watch	3.09	0.29
251	watering can	2.58	-0.04
252	watermelon	3.16	0.34
253	well	2.66	0.01
254	wheel	3.55	0.60
255	whistle	3.06	0.27
256	windmill	1.44	-0.80
257	window	0.13	-1.67
258	wineglass	3.24	0.39
259	wrench	1.37	-0.84
260	zebra	3.32	0.45




parts intact, no added parts



parts intact, added parts



parts deleted, no added parts



parts deleted, added parts

Deleting versus Adding Parts






Figure 2
(Results, Exp 1)





0% degradation, 0° rotation



30% degradation, 45° rotation



30% degradation, 135° rotation



60% degradation, 45° rotation



60% degradation, 135° rotation







Figure 5
(continued, exp 4)




0% degradation, no noise




30% degradation, sparse noise



30% degradation, dense noise



60% degradation, sparse noise



60% degradation, dense noise

Noise Density versus Degradation




Figure 7
(Results, Exp 5)